arXiv:1604.07153v1 [cs.DS] 25 Apr 2016

Closing the Gap for Makespan Scheduling via Sparsification Techniques*

Klaus Jansen[1], Kim-Manuel Klein[1], and José Verschae[2]

[1]University of Kiel, Department of Computer Science, Kiel, Germany.
E-mail:{kj,kmk}@informatik.uni-kiel.de
[2]Pontificia Universidad Católica de Chile, Facultad de Matemáticas & Escuela de Ingeniería, Santiago, Chile. E-mail:jverschae@uc.cl

April 26, 2016

Abstract

Makespan scheduling on identical machines is one of the most basic and fundamental packing problems studied in the discrete optimization literature. It asks for an assignment of n jobs to a set of m identical machines that minimizes the makespan. The problem is strongly NP-hard, and thus we do not expect a $(1+\varepsilon)$-approximation algorithm with a running time that depends polynomially on $1/\varepsilon$. Furthermore, Chen et al. [3] recently showed that a running time of $2^{(1/\varepsilon)^{1-\delta}} + \text{poly}(n)$ for any $\delta > 0$ would imply that the Exponential Time Hypothesis (ETH) fails. A long sequence of algorithms have been developed that try to obtain low dependencies on $1/\varepsilon$, the better of which achieves a running time of $2^{\tilde{O}(1/\varepsilon^2)} + O(n \log n)$ [11]. In this paper we obtain an algorithm with a running time of $2^{\tilde{O}(1/\varepsilon)} + O(n \log n)$, which is tight under ETH up to logarithmic factors on the exponent.

Our main technical contribution is a new structural result on the *configuration-IP*. More precisely, we show the existence of a highly symmetric and sparse optimal solution, in which all but a constant number of machines are assigned a configuration with small support. This structure can then be exploited by integer programming techniques and enumeration. We believe that our structural result is of independent interest and should find applications to other settings. In particular, we show how the structure can be applied to the minimum makespan problem on related machines and to a larger class of objective functions on parallel machines. For all these cases we obtain an efficient PTAS with running time $2^{\tilde{O}(1/\varepsilon)} + \text{poly}(n)$.

1 Introduction

Minimum makespan scheduling is one of the foundational problems in the literature on approximation algorithms [7, 8]. In the *identical machine* setting the problem asks for an assignment of a set of n jobs $\mathcal{J}$ to a set of m identical machines $\mathcal{M}$. Each job $j \in \mathcal{J}$ is characterized by a non-negative processing time $p_j \in \mathbb{Z}_{>0}$. The load of a machine is the total processing time of jobs assigned to it, and our objective is to minimize the *makespan*, that is, the maximum machine load. This problem is usually denoted $P||C_{\max}$. It is well known to admit a *polynomial time approximation scheme* (PTAS) [10], and there has been many subsequent works improving the running time or deriving PTAS's for more general settings. The fastest PTAS for $P||C_{\max}$ achieves a running time of $2^{O(1/\varepsilon^2)\log^3(1/\varepsilon))} + O(n \log n)$ for $(1+\varepsilon)$-approximate solutions [11]. Very recently, Chen et al. [3] showed that, assuming the *exponential time hypothesis*

*This work was partially supported by DFG Project, Entwicklung und Analyse von effizienten polynomiellen Approximationsschemata für Scheduling- und verwandte Optimierungsprobleme, Ja 612/14-2, by FONDECYT project 3130407, and by Nucleo Milenio Información y Coordinación en Redes ICM/FIC RC130003.

(ETH), there is no PTAS that yields $(1+\varepsilon)$-approximate solutions for $\varepsilon > 0$ with running time $2^{(1/\varepsilon)^{1-\delta}} + \text{poly}(n)$ for any $\delta > 0$ [3].

Given a guess $T \in \mathbb{N}$ on the optimal makespan, which can be found with binary search, the problem reduces to deciding the existence of a packing of the jobs to m machines (or bins) of capacity T. If we aim for a $(1+\varepsilon)$-approximate solution, for some $\varepsilon > 0$, we can assume that all processing times are integral and T is a constant number, namely $T \in O(1/\varepsilon^2)$. This can be achieved with well known rounding and scaling techniques [1, 2, 9] which will be specified later. Let $\pi_1 < \pi_2 < \ldots < \pi_d$ be the job sizes appearing in the instance after rounding, and let b_k denote the number of jobs of size π_k. The mentioned rounding procedure implies that the number of different job sizes is $d = O((1/\varepsilon)\log(1/\varepsilon))$. Hence, for large n we obtain a highly symmetric problem where several jobs will have the same processing time. Consider the *knapsack polytope* $\mathcal{P} = \{c \in \mathbb{R}^d_{\geq 0} : \pi \cdot c \leq T\}$. A packing on one machine can be expressed as a vector $c \in Q = \mathbb{Z}^d \cap \mathcal{P}$, where c_k denotes the number of jobs of size π_k assigned to the machine. Elements in $Q = \mathbb{Z}^d \cap \mathcal{P}$ are called *configurations*. Considering a variable $x_c \in \mathbb{Z}_{\geq 0}$ that decides the multiplicity of configuration c in the solution, our problem reduces to solving the following linear integer program (ILP):

$$\text{[conf-IP]} \quad \sum_{c \in Q} c \cdot x_c = b, \tag{1}$$

$$\sum_{c \in Q} x_c = m, \tag{2}$$

$$x_c \in \mathbb{Z}_{\geq 0} \qquad \text{for all } c \in Q. \tag{3}$$

In this article we derive new insights on this ILP that help us to design faster algorithms for $P||C_{\max}$ and other more general problems. These including makespan scheduling on *related machines* $Q||C_{\max}$, and a more general class of objective functions on parallel machines. We show that all these problems admit a PTAS with running time $2^{O((1/\varepsilon)\log^4(1/\varepsilon))} + \text{poly}(n)$. Hence, our algorithm is best possible up to polylogarithmic factors in the exponent assuming ETH [3].

1.1 Literature Review

There is an old chain of approximation algorithms for $P||C_{\max}$, starting from the seminal work by Graham [7, 8]. The first PTAS was given by Hochbaum and Shmoys [10] and had a running time of $(n/\varepsilon)^{O((1/\varepsilon)^2)} = n^{O((1/\varepsilon)^2 \log(1/\varepsilon))}$. This was improved to $n^{O((1/\varepsilon)\log^2(1/\varepsilon))}$ by Leung [16]. Subsequent articles improve further the running time. In particular Hochbaum and Shmoys (see [9]) and Alon et al. [1,2] obtain an *efficient PTAS*[1] (EPTAS) with running time $2^{(1/\varepsilon)^{\text{poly}(1/\varepsilon)}} + O(n \log n)$. Alon et al. [1,2] consider general techniques that work for several objective functions, including all L_p-norm of the loads and maximizing the minimum machine load.

The fastest PTAS known up to date for $P||C_{\max}$ achieves a running time of $2^{O((1/\varepsilon)^2 \log^3(1/\varepsilon))} + O(n \log n)$ [11]. More generally, this work gives an EPTAS for the case of related (uniform) machines, where each machine $i \in \mathcal{M}$ has a speed s_i and assigning to i job j implies a processing time of p_j/s_i. For this more general case the running time is $2^{O((1/\varepsilon)^2 \log^3(1/\varepsilon))} + \text{poly}(n)$. For the simpler case of $P||C_{\max}$, the ILP can be solved directly since the number of variables is a constant. This can be done with Lentras' algorithm [15], or even with Kannan's algorithm [14] that gives an improved running time. This technique yields a running time that is doubly exponential in $1/\varepsilon$. This was, in essence, the approach by Alon et al. [1,2] and Hochbaum and Shmoys [9]. To lower the dependency on $1/\varepsilon$, Jansen [11] uses a result by Eisenbrand and Shmonin [4] that implies the existence of a solution x with support of size at most $O(d \log(dT)) = O((1/\varepsilon)\log^2(1/\varepsilon))$. First guessing the support and then solving the ILP with $O((1/\varepsilon)\log^2(1/\varepsilon))$ integer variables and using Kannan's algorithm yields the desired running time of $2^{O((1/\varepsilon)^2 \log^3(1/\varepsilon))} + O(n \log n)$.

[1] That is, a PTAS whose running time is $f(1/\varepsilon)\text{poly}(|I|)$ where $|I|$ is the encoding size of the input and f is some function.

The configuration ILP has recently been studied in the context of the (1-dimensional) cutting stock problem. In this case, the dimension d is constant, $T = 1$, and π is a rational vector. Moreover, π and b are part of the input. Goemans and Rothvoß [5] obtain an optimal solution in time $\log(\Delta)^{2^{O(d)}}$, where Δ is the largest number appearing in the denominator of π_k or the multiplicities b_k. This is achieved by first showing that there exists a pre-computable set $\tilde{Q} \subseteq Q$ with polynomial many elements, such that there exists a solution x that gives all but constant (depending only on d) amount of weight to $\tilde{Q}$. We remark that applying this result to a rounded instance of $P||C_{\max}$ yields a running time that is doubly exponential on $1/\varepsilon$.

1.2 Our Contributions

Our main contribution is a new insight on the structure of the solutions of [conf-IP]. These properties are specially tailored to problems in which T is bounded by a constant, which in the case of $P||C_{\max}$ can be guaranteed by rounding and scaling. The same holds for $Q||C_{\max}$ with a more complex rounding and case analysis.

We first classify configurations by their support. We say that a configuration is *simple* if its support is of size at most $\log(T+1)$, otherwise it is *complex*. Our main structural result[2] states that there exists a solution x in which all but $O(d\log(dT))$ weight is given to simple configurations, the support is bounded by $O(d\log(dT))$ (as implied by Eisenbrand and Shmonin [4]) and no complex configuration has weight larger than 1.

Theorem 1 (Thin solutions)**.** *Assume that [conf-IP] is feasible. Then there exists a feasible solution x to [conf-IP] such that:*

1. *if $x_c > 1$ then the configuration c is simple,*
2. *the support of x satisfies $|\operatorname{supp}(x)| \leq 4(d+1)\log(4(d+1)T)$, and*
3. *$\sum_{c\in Q_c} x_c \leq 2(d+1)\log(4(d+1)T)$, where Q_c denotes the set of complex configurations.*

We call a solution satisfying the properties of the theorem *thin*. The theorem can be shown by iteratively applying a sparsification lemma that shows that if a solution gives a weight of two or more to a complex configuration, then we can replace this partial solution by two configurations with smaller support. The sparsification lemma is shown by a simple application of the pigeonhole principle. The theorem can be shown by mixing this technique with the theorem of Eisenbrand and Shmonin [4] and a potential function argument.

As an application to our main structural theorem, we derive a PTAS for $P||C_{\max}$ by first guessing the jobs assigned to complex configurations. An optimal solution for this subinstance can be derived by a dynamic program. For the remaining instance we know the existence of a solution using only simple configurations. Then we can guess the support of such solution and solve the corresponding [conf-IP] restricted to the guessed variables. The main use of having simple configurations is that we can guess the support of the solution much faster, as the number of simple configuration is (asymptotically) smaller than the total number of configurations. The complete procedure takes time $2^{O((1/\varepsilon)\log^4(1/\varepsilon))} + O(n\log n)$. Moreover, using the rounding and case analysis of Jansen [11], we derive an mixed integer linear program that can be suitably decomposed in order to apply our structural result iteratively. This yields a PTAS with a running time of $2^{O((1/\varepsilon)\log^4(1/\varepsilon))} + \operatorname{poly}(n)$ for $Q||C_{\max}$.

Similarly, we can extend our results to derive PTAS's for a larger family of objective functions as considered by Alon et al. [1,2]. Let ℓ_i denote the load of machine i, that is, the total processing time of jobs assigned to machine i for a given solution. Our techniques then gives a PTAS with the same running time for the problem of minimizing the L_p-norms of the loads (for fixed p), and maximizing $\min_{i\in M}\ell_i$, among others. To solve this problem, we can round the instance and

[2] We remark the resemblance of this structure to the result by Goemans and Rothvoß [5]. Indeed, similar to their result, we can precompute a subset of configurations such that all but a constant amount of weight of the solution is given to such set. In their case the set is of cardinality polynomial on the input and is constructed by covering the integral solutions of the knapsack polytope by parallelepipeds. In our case, all but $O(d\log dT)$ weight is given to simple configurations.

state an IP analogous to [conf-IP] but considering an objective function. However, the objective function prevents us to use the main theorem as it is stated. To get over this issue, we study several ILPs. In each ILP we consider x_c to be a variable only if c has a given load, and fix the rest to be some optimal solution. Applying to each such ILP Theorem 1, plus some extra ideas, yields an analogous structural theorem. Afterwards, an algorithm similar to the one for makespan minimization yields the desired PTAS.

From an structural point of view, our sparsification lemma has other consequences on the structure of the knapsack polytope and the LP-relaxation of the [conf-IP]. More precisely, we can show that any vertex of the convex hull of Q must be simple. This, for example, helps us to upper bound the number of vertices by $2^{O(\log^2(T)+\log^2(d))}$. Moreover, we can show that the configuration-LP, obtained by replacing the integrality restriction in [conf-IP] by $x \geq 0$, if it is feasible then admits a solution whose support consist purely of simple configurations. Due to space limitations we leave many details and proofs to the appendix.

2 Preliminaries

We will use the following notation throughout the paper. By default $\log(\cdot) = \log_2(\cdot)$, unless stated otherwise. Given two sets A, I, we will denote by A^I the set of all vectors indexed by I with entries in A, that is, $A^I = \{(a_i)_{i\in I} : a_i \in A \text{ for all } i \in I\}$. Moreover, for $A \subseteq \mathbb{R}$, we denote the support of a vector $a \in A^I$ as $\operatorname{supp}(a) = \{i \in I : a_i \neq 0\}$.

We consider an arbitrary knapsack polytope $\mathcal{P} = \{c \in \mathbb{R}^d_{\geq 0} : \pi \cdot c \leq T\}$ where $\pi \in \mathbb{Z}^d_{>0}$ is a non-negative integral (row) vector and T is a positive integer. We assume without loss of generality that each coordinate π_k of π is upper bounded by T (otherwise $c_k = 0$ for all $c \in \mathbb{Z}^d \cap \mathcal{P}$). We focus on the set of integral vectors in $\mathcal{P}$ which we denote by $Q = \mathbb{Z}^d \cap \mathcal{P}$. We call an element $c \in Q$ a *configuration*. Given $b \in \mathbb{R}^d$, consider the problem of decomposing b as a conic integral combination of m configurations. That is, our aim is to find a feasible solution to [conf-IP], defined above.

A crucial property of the [conf-IP] is that there is always a solution with a support of small cardinality. This follows from a Caratheodory-type bound obtained by Eisenbrand and Shmonin [4]. Since we will need the argument later, we state the result applied to our case and revise its (very elegant) proof. We split the proof in two lemmas.

For a given subset $A \subseteq Q$, let us denote by x^A the indicator vector of A, that is $x^A_c = 1$ if $c \in A$, and 0 otherwise. Let us also denote by M the $(d+1) \times |Q|$ matrix that defines the system of equalities (1) and (2).

Lemma 2 (Eisenbrand and Shmonin [4])**.** *Let $x \in \mathbb{Z}^Q_{\geq 0}$ be a vector such that $|\operatorname{supp}(x)| > 2(d+1)\log(4(d+1)T)$. Then there exist two disjoint sets A, B with $\emptyset \neq A, B \subseteq \operatorname{supp}(x)$ such that $Mx^A = Mx^B$.*

Proof. Let $s := |\operatorname{supp}(x)|$. Each coordinate of M is smaller than T. Hence, for any $A \subseteq \operatorname{supp}(x)$, each coordinate of Mx^A is no larger than $|A| \cdot T \leq sT$. Thus, Mx^A belongs to $\{0, \ldots, sT\}^{d+1}$, and hence there are at most $(sT+1)^{d+1} = 2^{(d+1)\log(sT+1)}$ different possibilities for vector Mx^A, over all possible subsets $A \subseteq \operatorname{supp}(x)$. On the other hand, there are 2^s many different subsets of $\operatorname{supp}(x)$.

We claim that $s > (d+1)\log(sT+1)$. Indeed, since $s > 2(d+1)\log(4(d+1)T)$ then $T < 2^{\frac{s}{2(d+1)}}/(4(d+1))$. Hence,

$$
\begin{aligned}
(d+1)\log(sT+1) &< (d+1)\log\left(\frac{s2^{\frac{s}{2(d+1)}}}{4(d+1)}+1\right)\\
&\leq (d+1)\log\left(2^{\frac{s}{2(d+1)}}\left(\frac{s}{4(d+1)}+1\right)\right)\\
&= (d+1)\left(\frac{s}{2(d+1)}+\log\left(\frac{s}{4(d+1)}+1\right)\right)\\
&\leq \frac{s}{2}+\frac{s}{4\ln(2)} < s,
\end{aligned}
$$

where the penultimate inequality follows since $\log(x) \leq (x-1)/\ln(2)$ for all $x \geq 1$.

We obtain that $2^s > 2^{(d+1)\log(sT+1)}$. Hence, by the pigeonhole principle there are two distinct subsets $A', B' \subseteq \operatorname{supp}(x)$ such that $Mx^{A'} = Mx^{B'}$. We can now define $A = A' \setminus B'$ and $B = B' \setminus A'$ and obtain $Mx^A = Mx^B$. It remains to show that $A, B \neq \emptyset$. Notice that if $A = \emptyset$ then $A' \subseteq B'$, and the last equality of $Mx^{A'} = Mx^{B'}$ implies that $|A'| = |B'|$. This is a contradiction since then $A' = B'$. We conclude that $A \neq \emptyset$. The proof that $B \neq \emptyset$ is analogous. □

Lemma 3 (Eisenbrand and Shmonin [4])**.** *If [conf-IP] is feasible, then there exists a feasible solution x such that $|\operatorname{supp}(x)| \leq 2(d+1)\log(4(d+1)T)$.*

Proof. Let x be a solution to [conf-IP] that minimizes $|\operatorname{supp}(x)| = s$. Assume by contradiction that $s > 2(d+1)\log(4(d+1)T)$. We show that we can find another solution x' to [conf-IP] with $|\operatorname{supp}(x')| < |\operatorname{supp}(x)|$, contradicting the minimality of $|\operatorname{supp}(x)|$. By Lemma 2, there exist two disjoint subsets $A, B \in \operatorname{supp}(x)$ such that $Mx^A = Mx^B$. Moreover, let $\lambda = \min\{x_c : c \in A\}$. Vector $x' := x - \lambda x^A + \lambda x^B$ is also a solution to [conf-IP] and has a strictly smaller support since a configuration $c^* \in \arg\min\{x_c : c \in A\}$ satisfies $x'_{c^*} = 0$. □

3 Structural Results

Recall that we call a configuration c simple if $|\operatorname{supp}(c)| \leq \log(T+1)$ and complex otherwise. An important observation to show Theorem 1 is that if c is a complex configuration, then $2c$ can be written as the sum of two configurations of smaller support. This is shown by the following Sparsification Lemma.

Lemma 4 (Sparsification Lemma)**.** *Let $c \in Q$ be a complex configuration. Then there exist two configurations $c_1, c_2 \in Q$ such that*

1. $\pi \cdot c_1 = \pi \cdot c_2 = \pi \cdot c$,
2. $2c = c_1 + c_2$,
3. $\operatorname{supp}(c_1) \subsetneq \operatorname{supp}(c)$ *and* $\operatorname{supp}(c_2) \subsetneq \operatorname{supp}(c)$.

Proof. Consider for each subset $S \subseteq \operatorname{supp}(c)$, a configuration $c^S \in Q$ such that $c^S_i = c_i$ if $i \in S$ and $c^S = 0$ otherwise. As the number of subsets of $\operatorname{supp}(c)$ is $2^{|\operatorname{supp}(c)|}$, and $c^R \neq c^S$ if and only if $R \neq S$, the collection of vectors $V := \{c^S : S \subseteq \operatorname{supp}(c)\}$ has cardinality $|V| = 2^{|\operatorname{supp}(c)|}$.

On the other hand, for any vector $c^S \in V$ it holds that $\pi \cdot c^S \leq \pi \cdot c \leq T$. Hence, $\pi \cdot c^S \in \{0, 1 \ldots, T\}$ can take only $T+1$ different values. Using that c is a complex configuration and hence $2^{|\operatorname{supp}(c)|} > 2^{\log(T+1)} = T+1$, the pigeonhole principle ensures that there are two different non-empty configurations $c^S, c^R \subseteq V$ with $\pi \cdot c^S = \pi \cdot c^R$. By removing the intersection, we can assume w.l.o.g. that S and R have no intersection. We define $c_1 = c - c^S + c^R$ and $c_2 = c - c^R + c^S$, which satisfy the properties of the lemma as

$$
\begin{aligned}
\pi \cdot c_1 &= \pi \cdot c - \pi \cdot c^S + \pi \cdot c^R = \pi \cdot c \quad \text{and}\\
2c &= c - c^S + c^R + c - c^R + c^S = c_1 + c_2.
\end{aligned}
$$

Since $\operatorname{supp}(c_1) \subseteq \operatorname{supp}(c) \setminus S$ and $\operatorname{supp}(c_2) \subseteq \operatorname{supp}(c) \setminus R$, property 3 is satisfied. □

With Lemma 4 we are ready to show Theorem 1. For the proof it is tempting to apply the lemma iteratively, replacing any complex configuration that is used twice by two configurations with smaller support. This can be repeated until there is no complex configuration taken multiple times. Then we can apply the technique of Lemma 3 to the obtained solution to bound the cardinality of the support. However, the last step might break the structure obtained if the solution implied by Lemma 3 uses a complex configuration more than once. In order to avoid this issue we consider a potential function. We show that a vector minimizing the chosen potential uses each complex configuration at most once, and that the number of complex configurations in the support is bounded. Finally, we apply the techniques from Lemma 3 restricted to variables corresponding to simple configurations.

Proof of Theorem 1. Consider the following potential function of a solution $x \in \mathbb{Z}^Q_{\geq 0}$ of [conf-IP],

$$\Phi(x) = \sum_{\text{complex config. } c} x_c |\operatorname{supp}(c)|.$$

Let x be a solution of [conf-IP] with minimum potential $\Phi(x)$, which is well defined since the set of feasible solutions has finite cardinality. We show two properties of x.

P1: $x_c \leq 1$ for each complex configuration $c \in Q$.

Assume otherwise. Consider the two configurations c_1 and c_2 implied by the previous lemma. We define a new solution $x'_e = x_e$ for $e \notin \{c, c_1, c_2\}$, $x'_{c_1} = x_{c_1} + 1$, $x'_{c_2} = x_{c_2} + 1$ and $x'_c = x_c - 2$. Since $|\operatorname{supp}(c_1)| < |\operatorname{supp}(c)|$ and $|\operatorname{supp}(c_2)| < |\operatorname{supp}(c)|$, we obtain that $\Phi(x') < \Phi(x)$ which contradicts the minimality of $\Phi(x)$.

P2: The number of complex configurations in $\operatorname{supp}(x)$ is at most $2(d+1)\log(4(d+1)T)$.

Let $\tilde{x}$ be the vector defined as $\tilde{x}_c = x_c$ if $c \in Q$ is complex, and $\tilde{x} = 0$ if $c \in Q$ is simple. Then Lemma 2 implies that there are exist two disjoint subsets $A, B \subseteq \operatorname{supp}(\tilde{x})$ of complex configurations such that $Mx^A = Mx^B$. Thus, the solution $x' = x - x^A + x^B$ and the solution $x'' = x - x^B + x^A$ are feasible for [config-IP]. By linearity, the potential function on the new solutions are $\Phi(x') = \Phi(x) - \Phi(x^A) + \Phi(x^B)$ or respectively $\Phi(x'') = \Phi(x) - \Phi(x^B) + \Phi(x^A)$. If $\Phi(x^A) > \Phi(x^B)$ or $\Phi(x^B) > \Phi(x^A)$ then we have constructed a new solution with smaller potential, contradicting our assumption on the minimality of $\Phi(x)$. We conclude that $\Phi(x^B) = \Phi(x^A)$ and thus $\Phi(x) = \Phi(x')$. By construction of x', we obtain that $x'_c > x_c \geq 1$ for any complex configuration $c \in B$. Having multiplicity ≥ 2 for a complex configuration c, we can proceed as in Case 1 to find a new solution with decreased potential, which yields a contradiction.

Given these two properties, to conclude the theorem it suffices to upper bound the number of simple configurations by $2(d+1)\log(4(d+1)T)$. Suppose this property is violated, then we find two sets $A, B \subseteq \operatorname{supp}(x)$ of simple configurations (see Lemma 2) with $Mx^A = Mx^B$ and proceed as in Lemma 3. Since Lemma 3 is only applied to simple configurations, properties P1 and P2 continue to hold and the theorem follows. □

Our techniques, in particular our Sparsification Lemma, imply two corollaries on the structure of the knapsack polytope and the LP-relaxation implied by the [conf-IP].

Corollary 5. *Every vertex of conv.hull(Q) is a simple configuration. Moreover, the total number of simple configurations in Q is upper bounded by $2^{O(\log^2(T)+\log^2(d))}$ and thus the same expression upper bounds the number of vertices of conv.hull(Q).*

Proof. Consider a complex configuration $c \in Q$. By Lemma 4 we know that there exist $c_1, c_2 \in Q$ with $c_1, c_2 \neq c$ such that $2c = c_1 + c_2$. Hence, c is not a vertex of Q as it can be written as a convex combination $c = c_1/2 + c_2/2$.

To bound the number of simple configurations fix a set $D \subseteq \{1, \ldots, d\}$. Notice that the number of configurations c with $\operatorname{supp}(c) = D$ is at most $T^{|D|}$. For simple configurations it suffices to take D with cardinality at most $\log(T+1)$. Since the number of subsets $D \subseteq \{1, \ldots, d\}$ with

cardinality i is $\binom{d}{i}$, we obtain that the number of simple configurations is at most

$$\sum_{i=0}^{\lfloor \log(T+1) \rfloor} \binom{d}{i} \times (T+1)^{\log(T+1)} \leq (\log(T+1)+1) d^{\log(T+1)} \times (T+1)^{\log(T+1)}$$
$$= 2^{\log(\log(T+1)+1)+\log(d)\log(T+1)} \times 2^{\log(T+1)\log(T+1)} = 2^{O(\log^2(d)+\log^2(T))}. \qquad \square$$

The following corollary follows as each complex configuration can be represented by a convex combination of simple configurations.

Corollary 6. *Let [conf-LP] be the LP relaxation of [conf-IP], obtained by changing each constraint $x_c \in \mathbb{Z}_{\geq 0}$ to $x_c \geq 0$ for all $c \in Q$. If the LP is feasible then there exists a solution x such that each configuration $c \in \operatorname{supp}(x)$ is simple.*

Proof. Consider a solution x of [conf-LP]. Assume that there exists $c \in Q$ such that c is complex and $x_c > 0$. Then by the previous corollary, configuration c can be written as $c = \sum_{q \in Q} \lambda_q q$, where $\sum_{q \in Q} \lambda_q = 1$, $\lambda_q \geq 0$ for all $q \in Q$, and $\lambda_q = 0$ if $q \in Q$ is complex. Consider a new solution x' defined as

$$x'_q = \begin{cases} 0 & \text{if } q = c, \\ x_q + \lambda_q \cdot x_c & \text{if } q \neq c. \end{cases}$$

This new solution is also feasible for [conf-LP]. As $x'_c = 0$, the number of complex configurations in the support of the solution is reduced by 1. This procedure can be repeated until we have a solution $\hat{x}$ whose support contains only simple configurations. $\square$

4 Applications to Scheduling on Parallel Machines

In what follows we show how to exploit the structural insights of the previous section to derive faster algorithms for parallel machines scheduling problems. We start by considering $P||C_{\max}$, where we seek to assign a set of jobs $\mathcal{J}$ with processing times $p_j \in \mathbb{Z}_{>0}$ to a set $\mathcal{M}$ of m machines. For a given assignment $a : \mathcal{J} \mapsto \mathcal{M}$, we define the load of a machine i as $\sum_{j:a(j)=i} p_j$ and the *makespan* as the maximum load of jobs over all machines, which is the minimum time needed to complete the execution of all jobs on the processors. The goal is to find an assignment $\mathcal{J} \mapsto \mathcal{M}$ that minimizes the makespan.

We first follow well known rounding techniques [1, 2, 9, 10]. Consider an error tolerance $0 < \varepsilon < 1/3$ such that $1/\varepsilon^2$ is an integer. To get an estimation of the optimal makespan, we follow the standard dual approximation approach. First, we can use, e.g., the 2-approximation algorithm by Graham [7] to get an initial guess of the optimal makespan. Using binary search, we can then estimate the optimal makespan within a factor of $(1+\varepsilon)$ in $O(\log(1/\varepsilon))$ iterations. Therefore, it remains to give an algorithm that decides for a given makespan T, if there exists an assignment with makespan $(1+O(\varepsilon))T$ or reports that there exists no assignment with makespan $\leq T$.

For a given makespan T we define the set of big jobs $\mathcal{J}_{\text{big}} = \{j \in \mathcal{J} : p_j \geq \varepsilon T\}$ and the set of small jobs $\mathcal{J}_{\text{small}} = \mathcal{J} \setminus \mathcal{J}_{\text{big}}$. The following lemma shows that small jobs can be replaced from the instance by adding big jobs, each of size εT, as placeholders. Let S be the sum of processing times of jobs in $\mathcal{J}_{\text{small}}$ and let S^* denote the next value of S rounded up to the next multiple of εT, that is, $S^* = \varepsilon T \cdot \lceil S/(\varepsilon T) \rceil$. We define a new instance containing only big jobs by $\mathcal{J}^* = \mathcal{J}_{\text{big}} \cup \mathcal{J}_{\text{new}}$, where $\mathcal{J}_{\text{new}}$ contains $S^*/(\varepsilon T) \in \mathbb{N}$ jobs of size εT.

Lemma 7. *Given a feasible assignment $a : \mathcal{J} \mapsto \mathcal{M}$ of jobs with makespan T. Then there exists a feasible assignment $a_B : \mathcal{J}^* \mapsto \mathcal{M}$ of makespan $T^* \leq (1+\varepsilon)T$. Similarly, an assignment of jobs in $\mathcal{J}^*$ of makespan T^* can be transformed to an assignment of $\mathcal{J}$ of makespan at most $(1+\varepsilon)T^*$.*

Proof. We modify the assignment a of jobs in $\mathcal{J}$ by replacing the set of small jobs on each machine by jobs in $\mathcal{J}_{\text{new}}$. Let S_i be the total processing time of small jobs assigned to machine i. Then the small jobs are replaced by (at most) $S_i^*/(\varepsilon T)$ jobs in $\mathcal{J}_{\text{new}}$, where S_i^* denotes the

value of S_i rounded up to the next multiple of εT. As $\sum \frac{S_i^*}{\varepsilon T} \geq \lfloor \sum \frac{S_i}{\varepsilon T} \rfloor = \frac{S^*}{\varepsilon T}$, the new solution processes all jobs in $\mathcal{J}_{\text{new}}$ and the load on each machine increases hence by at most εT. Having an assignment for the big jobs $\mathcal{J}^*$, we can easily obtain a schedule for jobs $\mathcal{J}$, by adding the small items greedily into the space of the placeholder jobs $\mathcal{J}_{\text{new}}$. □

By scaling the processing times of jobs in $\mathcal{J}^*$, we can assume that the makespan T has value $1/\varepsilon^2$. Also notice that we can assume that $p_j \leq T$ for all j, otherwise we cannot pack all jobs within makespan T. This implies that each job $j \in \mathcal{J}^*$ has a processing time of $1/\varepsilon \leq p_j \leq 1/\varepsilon^2$. In the following we give a transformation of big jobs in $\mathcal{J}^*$ by rounding their processing times. We first round the jobs to the next power of $1+\varepsilon$ as $p'_j = (1+\varepsilon)^{\lceil \log_{(1+\varepsilon)} p_j \rceil}$, and thus all rounded processing times belong to $\Pi' = \{(1+\varepsilon)^k : 1/\varepsilon \leq (1+\varepsilon)^k \leq (1+\varepsilon)/\varepsilon^2 \text{ and } k \in \mathbb{N}\}$. We further round processing times p'_j to the next integer $\bar{p}_j = \lceil p'_j \rceil$ and define a new set $\Pi = \{\lceil p \rceil : p \in \Pi'\}$. Notice that Π only contains integers and $|\Pi| \leq |\Pi'| \in O((1/\varepsilon)\log(1/\varepsilon))$.

Lemma 8. *If there is a feasible schedule of jobs $\mathcal{J}^*$ with processing times p_j onto m machines with makespan $T^* \leq (1+\varepsilon)T$, then there is also a feasible schedule of jobs $\mathcal{J}^*$ with rounded processing $\bar{p}_j$ with a makespan of at most $(1+5\varepsilon)T$. Furthermore, the number of different processing times is at most $|\Pi| \in O((1/\varepsilon)\log(1/\varepsilon))$.*

Proof. Consider a feasible schedule of jobs in $\mathcal{J}^*$ with processing times p_j onto m machines with makespan T^*. Let $J_{i_1}, \ldots, J_{i_r}$ be the set of jobs processed on machine i i.e. $a(J_{i_k}) = i$ for $k = 1, \ldots, r$. Then $\sum_{j=1}^{r} p'_j \leq \sum_{j=1}^{r}(1+\varepsilon)p_j \leq (1+\varepsilon)T^*$. Hence, the same assignment a with processing times p'_j yields a makespan of at most $(1+\varepsilon)T^* \leq (1+\varepsilon)^2 T = 1/\varepsilon^2 + 2/\varepsilon + 1$. Since $p'_j \geq p_j \geq 1/\varepsilon$, on every machine are at most $1/\varepsilon + 2$ jobs. Hence, rounding the processing times p'_j of each job to the next integer increase the load on each machine by at most $1/\varepsilon + 2$. Recalling that $\varepsilon < 1/3$, we obtain a feasible schedule with makespan at most $(1+\varepsilon)T^* + 1/\varepsilon + 2 \leq 1/\varepsilon^2 + 3/\varepsilon + 3 < T + 5\varepsilon T$. □

In what follows we give an algorithm that decides in polynomial time the existence of a solution for instance $\mathcal{J}^*$ with processing times $\bar{p}_j$ and makespan $\bar{T} = \lfloor (1+5\varepsilon)T \rfloor$. We call numbers in Π by $\pi_1, \ldots, \pi_d$ and define the vector $\pi = (\pi_1, \pi_2, \ldots, \pi_d) \in \mathbb{N}^d$ of rounded processing times. We consider *configurations* to be vectors in $Q = \mathcal{P} \cap \mathbb{Z}^d$, where $\mathcal{P} = \{c \in \mathbb{R}^d_{\geq 0} : \pi \cdot c \leq \bar{T}\}$ is a knapsack polytope (see Section 3). As before, we say that a configuration is simple if $|\operatorname{supp}(c)| \leq \log(\bar{T}+1)$, and complex otherwise. For a given assignment of jobs to machines, we say that a machine follows a configuration c if c_k is the number of jobs of size π_k assigned to the machine. We denote by $Q_c \subseteq Q$ the set of complex configurations and by $Q_s \subseteq Q$ the set of simple configurations.

Let b_k be the number of jobs of size π_k in the instance $\mathcal{J}^*$ (with processing times $\bar{p}$). Consider an ILP with integer variables x_c for each $c \in Q$, which denote the number of machines that follow configuration c. With these parameters the problem of scheduling all jobs in a solution of makespan $\bar{T}$ is equivalent to finding a solution to [conf-IP]. To solve the ILP we use, among other techniques, Kannan's algorithm [14] which is an improvement on the algorithm by Lenstra [15]. The algorithm has a running time of $2^{O(N \log N)} s$ where N is the number of variables and s is number of bits used to encode the input of the ILP in binary.

By Theorem 1, if [conf-IP] is feasible then there exists a thin solution. In particular if one configuration c is used by more than one machine then c is simple, and the total number of used configurations is $4(d+1)\log(4(d+1)\bar{T}) \in O((1/\varepsilon)\log^2(1/\varepsilon))$. Additionally, the number of machines following a complex configurations is at most $2(d+1)\log(4(d+1)\bar{T}) \in O((1/\varepsilon)\log^2(1/\varepsilon))$. We consider the following strategy to decide the existence of a schedule of makespan $\bar{T}$.

Algorithm 9.

1. *For each processing time π_k, guess the number $b_k^c \leq b_k$ of jobs covered by complex configurations.*

2. *Find a minimum number of machines m^c to schedule jobs b^c with makespan $\bar{T}$.*

3. *Guess the support of simple configurations $\bar{Q}_s \subseteq Q_s$ used by a thin solution, with $|\bar{Q}_s| \leq 4(d+1)\log(4(d+1)\bar{T}) \in O((1/\varepsilon)\log^2(1/\varepsilon))$.*

4. *Solve the ILP restricted to configurations in $\bar{Q}_s$:*

$$\begin{aligned} \sum_{c \in \bar{Q}_s} c \cdot x_c &= b - b^c, \\ \sum_{c \in \bar{Q}_s} x_c &= m - m^c, \\ x_c &\in \mathbb{Z}_{\geq 0} \qquad \text{for all } c \in \bar{Q}_s. \end{aligned}$$

One of the key observations to prove the running time of the algorithm is that the number of simple configurations $|Q_s|$ is bounded by a quasi polynomial term:

$$|Q_s| \leq 2^{O(\log^2(1/\varepsilon))}.$$

This follows easily by Corollary 5, using that $|\bar{T}| \in O(1/\varepsilon^2)$ and $d = |\Pi| \in O((1/\varepsilon)\log(1/\varepsilon))$.

Lemma 10. *Algorithm 9 can be implemented with a running time of $2^{O((1/\varepsilon)\log^4(1/\varepsilon))}\log(n)$.*

Proof. In step 1, the algorithm guesses which jobs are processed on machines following a complex configurations. Since each configuration contains at most $O(1/\varepsilon)$ jobs, there are at most $O(m^c/\varepsilon) = O((1/\varepsilon^2)\log^2(1/\varepsilon))$ jobs assigned to such machines. For each size $\pi_k \in \Pi$, we guess the number b_k^c of jobs of size π_k assigned to such machines. Hence, we can enumerate all possibilities for jobs assigned to complex machines in time $2^{O((1/\varepsilon)\log^2(1/\varepsilon))}$. After guessing the jobs, we can assign them to a minimum number of machines in step 2 (with makespan $\bar{T}$) with a simple dynamic program that stores vectors $(\ell, z_1, \ldots, z_d)$ with $z_k \leq b_k^c$ being the number of jobs of size π_k used in the first $\ell \leq m^c$ processors [13]. The size of the dynamic programming table is $O(m^c \prod_{k=1}^{d}(b_k^c+1))$. For any vector $(\ell, z_1, \ldots, z_d)$, determining whether it corresponds to a feasible solution can be done by checking all vectors of the type $(\ell-1, z_1', \ldots, z_d')$ for $z_k' \leq z_k$. Thus, the running time of the dynamic program is $O(m^c[\prod_{k=1}^{d}(b_k^c+1)]^2)$. Since $b_k^c \in O((1/\varepsilon^2)\log^2(1/\varepsilon))$ for each k, recalling that $m^c \in O((1/\varepsilon)\log^2(1/\varepsilon))$, and that $d = |\Pi| \in O((1/\varepsilon)\log(1/\varepsilon))$, we obtain that step 2 can be implemented with $2^{O((1/\varepsilon)\log^2(1/\varepsilon))}$ running time.

In step 3, our algorithm guesses the support of a thin solution x. Recall that if x is thin then $|\operatorname{supp}(x)| \leq 4(d+1)\log(4(d+1)\bar{T}) = O((1/\varepsilon)\log^2(1/\varepsilon))$. Let $D = 4(d+1)\log(4(d+1)\bar{T})$. Then this guess can be done in time

$$\sum_{i=0}^{D} \binom{|Q_s|}{i} \leq (D+1)|Q_s|^D \leq 2^{O((1/\varepsilon)\log^4(1/\varepsilon))}.$$

We remark that for this step is that thin solutions are particularly useful. Indeed, guessing the support on the original ILP takes time $2^{O((1/\varepsilon)^2\log^3(1/\varepsilon))}$.

In step 4, the number of variables of the restricted ILP is $4(d+1)\log(4(d+1)\bar{T}) = O((1/\varepsilon)\log^2(1/\varepsilon))$. Moreover, the size of the input can be bounded by $O((1/\varepsilon^2)\log^3(1/\varepsilon)\log(n))$. Running Kannan's algorithm [14] to solve the ILP takes time $2^{O((1/\varepsilon)\log^3(1/\varepsilon))}\log(n)$. Hence, the total running time of our algorithm can be bounded by $2^{O((1/\varepsilon)\log^4(1/\varepsilon))}\log(n)$. □

Putting all pieces together, we conclude with the following theorem.

Theorem 11. *The minimum makespan problem on parallel machines $P||C_{\max}$ admits an EPTAS with running time $2^{O((1/\varepsilon)\log^4(1/\varepsilon))} + O(n\log n)$.*

Proof. Consider a scheduling instance with job set $\mathcal{J}$, processing times p_j for $j \in \mathcal{J}$ and machine set $\mathcal{M}$. The greedy algorithm by Graham to obtain a 2-approximation can be implemented in $O(n\log n)$. After guessing the makespan T, the processing times are sorted and rounded as described in Lemma 8. The rounding step can easily be implemented in $O(n)$ time.

Applying Algorithm 9 after the rounding needs, according to Theorem 10, a running time of $2^{O((1/\varepsilon)\log^4(1/\varepsilon))}\log(n)$ time. Since there are at most $O(\log(1/\varepsilon))$ many guessing rounds for the makespan, we obtain a total running time of $O(n\log n + \log(1/\varepsilon)\cdot n) + 2^{O((1/\varepsilon)\log^4(1/\varepsilon))}\log(n)$.

If $n \leq 2^{\frac{1}{\varepsilon}\log^4(\frac{1}{\varepsilon})}$ then the running time is upper bounded by $2^{O(\frac{1}{\varepsilon}\log^4(\frac{1}{\varepsilon}))}$, otherwise, the running time is at most $O(n\log n)$. In any case, the running time can be bounded by $2^{O(\frac{1}{\varepsilon}\log^4(\frac{1}{\varepsilon}))} + O(n\log n)$. □

4.1 Extension to other objectives

We now consider a more general family of objective functions defined by Alon et al. [1,2]. For a fixed function $f : \mathbb{R}_{\geq 0} \to \mathbb{R}_{\geq 0}$, we consider the following two objective functions:

$$\text{(I) } \min \sum_{i\in\mathcal{M}} f(\ell_i) \qquad \text{(II) } \min \max_{i\in\mathcal{M}} f(\ell_i),$$

where ℓ_i denotes the load of machine i. Analogously, we study maximization versions of the problems

$$\text{(I')} \max \sum_{i\in\mathcal{M}} f(\ell_i) \qquad \text{(II')} \max \min_{i\in\mathcal{M}} f(\ell_i),$$

For the minimization versions of the problem we assume that f is convex, while for (I') and (II') we assume it is concave. Moreover, we will need that the function satisfies the following sensitivity condition.

Condition 12. *For all $\varepsilon > 0$ there exists $\delta = \delta(\varepsilon) > 0$ such that for all $x, y \in \mathbb{R}_{\geq 0}$,*

$$(1-\delta)y \leq x \leq (1+\delta)y \quad \Rightarrow \quad (1-\varepsilon)f(y) \leq f(x) \leq (1+\varepsilon)f(y).$$

Alon et al. showed that each problem in that family admits a PTAS with running time $h(\varepsilon) + O(n\log n)$, where $h(\varepsilon)$ is a constant term that depends only on ε. Moreover, if $\delta(\varepsilon)$ in the condition further satisfies that $1/(\delta(\varepsilon)) \in O(1/\varepsilon)$, the running time is $2^{(1/\varepsilon)^{\mathrm{poly}(1/\varepsilon)}} + O(n\log n)$. In what follows we show how to improve this dependency. Since $1/(\delta(\varepsilon)) \in O(1/\varepsilon)$, we know that, for small enough ε, there exists a constant γ (independent of ε and δ) such that $1/\delta \leq \gamma/\varepsilon$. Moreover, we can assume w.l.o.g. that $\delta \leq \varepsilon$, and thus $\delta \leq \varepsilon \leq \gamma\delta$.

It is worth noticing that many interesting functions belong to this family. In particular (II) with $f(x) = x$ corresponds to the minimum makespan problem, (I) with $f(x) = x^p$, for constant p, corresponds to a problem that is equivalent to minimizing the L_p-norm of the vector of loads. Similarly, (II') with $f(x) = x$ corresponds to maximizing the minimum machine load. Notice that for all those objectives we have that $1/\delta = O(1/\varepsilon)$.

The techniques of Alon et al. [2] are based on a rounding method and then solving an ILP. We based our results in the same rounding techniques. Consider an arbitrary instance of a scheduling problem on identical machines with objective function (I), (II), (I') or (II'). Their first observation is that, if $L = \sum_j p_j/m$ is the average machine load, then a job with $p_j \geq L$ is scheduled alone on a machine in an optimal solution [2]. Hence, we can remove such job and a machine from the instance. In what follows, we assume without loss of generality, that $p_j < L$ for all j. For the sake of brevity, we summarize the rounding techniques of Alon et al. in the following theorem.

Theorem 13 (Alon et al. [2])**.** *Consider an instance for the scheduling problem with job set $\mathcal{J}$, identical machines $\mathcal{M}$, and processing times p_j for $j \in \mathcal{J}$ such that $p_j < L$ for all j. There exists a linear time algorithm that creates a new instance I' with job set $\mathcal{J}'$, machine set $\mathcal{M}$, and processing times p'_j. Moreover, there is an integer $\lambda \geq 1/\delta$ with $\lambda \in O(1/\delta)$ such that the new instance satisfies the following:*

1. *Each job j in I' has processing time $L/\lambda \leq p'_j \leq L$, and p'_j is a integer multiple of L/λ^2.*
2. *If $L' = \sum_j p'_j/m$ then $L \leq L' \leq (1+2/\lambda)L$.*
3. *Let $\textsc{Opt}$ and $\textsc{Opt}'$ be the optimal value of instances I and I', respectively. Then $(1-\varepsilon)\textsc{Opt} \leq \textsc{Opt}' \leq (1+\varepsilon)\textsc{Opt}$.*

4. *There exists a linear time algorithm that transforms a feasible solution for instance I' with objective value V to a feasible solution for I with objective value V' such that $(1-\varepsilon)V \leq V' \leq (1+\varepsilon)V$.*

Given this result, it suffices to find a $(1+\varepsilon)$-approximate solution for instance I'. To do so, we further round the processing times as in the previous section by defining $\bar{p}_j$ as the value $(1+\delta)^{\lceil \log_{1+\delta} p'_j \rceil}$ rounded up to the next multiple of L/λ^2 for all $j \in \mathcal{J}'$. Notice that $\bar{p}_j \leq (1+\delta)^{\lceil \log_{1+\delta} p'_j \rceil} + L/\lambda^2 \leq (1+\delta)p'_j + L/\lambda^2 \leq (1+\delta)p'_j + p'_j/\lambda \leq (1+2\delta)p'_j \leq (1+\delta)^2 p'_j$. Hence, for any assignment that gives a load ℓ_i on machine i for p'_j, the same assignment has a load $\bar{\ell}_i$ with $\ell_i \leq \bar{\ell}_i \leq (1+\delta)^2 \ell_i$. By Condition 12 we conclude that the new optimal value $\overline{\text{OPT}}$ satisfies that $(1-O(\varepsilon))\text{OPT} \leq \overline{\text{OPT}} \leq (1+O(\varepsilon))\text{OPT}$.

Let $\Pi = \{\pi_1, \ldots, \pi_d\}$ be the distinct values that the processing times $\bar{p}_j$ can take. Notice that $d = |\Pi| = O((1/\delta)\log(1/\delta))$. We consider the knapsack polytope with capacity $\bar{T} := 4L$, that is $\mathcal{P} = \{c \in \mathbb{R}^d_{\geq 0} : \pi \cdot c \leq \bar{T}\}$. Notice that π and $\bar{T}$ are integer multiples of L/λ^2, and thus $\mathcal{P} = \{c \in \mathbb{R}^d_{\geq 0} : \pi/(L/\lambda^2) \cdot c \leq \bar{T}/(L/\lambda^2)\}$. The following lemma, that is a simple adaptation of an observation by Alon et al. [2], shows that there exists an optimal solution for the rounded instance that uses only configurations in $\mathcal{P}$.

Lemma 14. *For $\varepsilon > 0$ small enough, the rounded instance with processing times $\bar{p}_j$ admits an optimal solution with makespan at most $4L$.*

Proof. Among all optimal solutions to the problem, consider one that minimizes $\sum_i \ell_i^2$, where ℓ_i is the load on machine i. Assume that there exists a machine i such that $\ell_i > 4L$. Notice that

$$\sum_j \bar{p}_j/m \leq (1+\delta)^2 \sum_j p'_j/m = (1+\delta)^2 L' \overset{\text{Theorem 13}}{\leq} (1+\delta)^2(1+2/\lambda)L \leq (1+\delta)^4 L.$$

Since $\delta \leq \varepsilon$, for ε small enough ($\varepsilon \leq 1/10$ suffices) we have that $(1+\delta)^4 \leq 2$ and thus $\sum_j \bar{p}_j/m \leq 2L$. Also, recall that $\bar{p}_j \leq (1+\delta)^2 p'_j \leq (1+\delta)^2 L \leq 2L$ for any j, where the second to last inequality follows from Theorem 13. Since $\ell_{\min} = \min_i \ell_i \leq \sum_j \bar{p}_j/m \leq 2L$, then $\ell_i - \ell_{\min} > 4L - 2L = 2L$. Then, for any job j, we have that $p_j < \ell_i - \ell_{\min}$. Let j^* be any job assigned to machine i. Hence, in particular we have that $p_{j^*} < \ell_i - \ell_{\min}$.

Recall that for problems (I) and (II) function f is convex. Hence, it holds that $f(x+\Delta) + f(y-\Delta) \leq f(x) + f(y)$ for all $0 \leq x \leq y$ with $0 \leq \Delta \leq y - x$ [2]. Moreover, the inequality becomes strict if f is strictly convex. Setting $x = \ell_{\min}$, $y = \ell_i$ and $\Delta = p_{j^*}$, the inequality implies that moving job j^* to machine $i^* \in \arg\min_i \ell_i$ decreases strictly $\sum_i \ell_i^2$. Moreover, the objective function (I) does not increase when performing this move, which yields a contradiction for this objective. Similarly, for problem (II) the objective does not increase since $\max\{f(\xi) : \xi \in [x, y]\}$ is always attained at x or y for f convex. This yields a contradiction for (II).

Analogously, for problems (I') and (II') function f is concave and thus $f(x+\Delta) + f(y-\Delta) \geq f(x) + f(y)$ holds for all $0 \leq x \leq y$ with $0 \leq \Delta \leq y - x$ [2]. Hence, moving job j^* to machine i^* decreases $\sum_i \ell_i^2$ but does not increase the objective (I'). Since f concave implies that $\min\{f(\xi) : \xi \in [x, y]\}$ is always attained at x or y, we also obtain a contradiction for (II'). The lemma follows. □

Let $L = \sum_j p_j/m$ be the average machine load (of the original instance). After our rounding we obtain an instance I' with job set $\mathcal{J}'$ and processing times $\bar{p}_j$ for $j \in \mathcal{J}'$. Moreover, the $\bar{p}_j$ are multiples of L/λ^2, where $\lambda \geq 1/\delta$ is an integer such that $\lambda = O(1/\delta)$, and also $\bar{p}_j \geq L/\lambda$. It holds that there exists an optimal solution of the rounded instance with makespan at most $4L$, see Lemma 14 (in particular $\bar{p}_j \leq 4L$ for all j). Let $\Pi = \{\pi_1, \ldots, \pi_d\}$ be the distinct values that the processing times $\bar{p}_j$ can take. Our rounding guarantees that $d = |\Pi| = O((1/\delta)\log(1/\delta))$. We consider the knapsack polytope with capacity $\bar{T} := 4L$, that is $\mathcal{P} = \{c \in \mathbb{R}^d_{\geq 0} : \pi \cdot c \leq \bar{T}\}$. Notice that π and $\bar{T}$ are integer multiples of L/λ^2, and that $\mathcal{P}$ can also be written as $\{c \in \mathbb{R}^d_{\geq 0} : \pi/(L/\lambda^2) \cdot c \leq \bar{T}/(L/\lambda^2)\}$.

As before, we say that a configuration is simple if $|\text{supp}(c)| \leq \log(\bar{T}+1)$, and complex otherwise. We denote by $Q_c \subseteq Q$ the set of complex configurations and by $Q_s \subseteq Q$ the set of simple configurations. In what follows we focus on objective function (I).

We set an ILP for the problem as before. Notice that each configuration c incurs a cost of $f_c := f(\pi \cdot c)$. Moreover, we round and scale the values f_c by defining $\bar{f}_c = \lceil f_c/(\varepsilon f_{\min}) \rceil$, where $f_{\min} = \min_{c \in Q} f_c$. It is not hard to see that solving a problem with those coefficients yields a $(1+\varepsilon)$-approximate solution to the optimal solution of I' with processing times $\bar{p}_j$. Let also b_k be the number of jobs j of processing time $\bar{p}_j = \pi_k$ in $\mathcal{J}'$. Consider the ILP obtained by adding to [conf-IP] the objective function $\min \sum_{c \in Q} \bar{f}_c \cdot x_c$. We call this ILP [cost-conf-IP]. With our previous discussion, it suffices to solve this ILP optimally. To solve this problem, we first notice that the largest coefficient in the objective can be bounded as follows.

Lemma 15. *If f satisfies Condition 12 then the largest value $\max_{c \in Q} \bar{f}_c$ is upper bounded by $1/\delta^{O(1)}$.*

Proof. We first bound $(\max_{c \in Q} f_c)/(\varepsilon f_{\min})$. Notice that Condition 12 implies that f is continuous on $\mathbb{R}_{\geq 0}$, and thus it admits a minimum and maximum in the interval $[L/\lambda, 4L]$. Let $x_{\min} \in \arg\min\{f(x) : x \in [L/\lambda, 4L]\}$ and $x_{\max} \in \arg\max\{f(x) : x \in [L/\lambda, 4L]\}$.

Consider first the case in which $x_{\min} \leq x_{\max}$ (this is not always true since f might not be monotone). We now use Condition 12 iteratively. Let $y^k := (1+\delta)^k x_{\min}$. Since $y^k \leq y^{k-1}(1+\delta)$, Condition 12 implies that $f(y^k) \leq (1+\varepsilon) f(y^{k-1})$. Iterating this idea we obtain that $f(y^k) \leq (1+\varepsilon)^k f(y^0)$. Taking $k = \lceil \log_{1+\delta}(x_{\max}/x_{\min}) \rceil$ implies that $x_{\max} \leq y^k \leq x_{\max}(1+\delta)$ and thus, by Condition 12, it holds that $f(y^k) \geq (1-\varepsilon) f(x_{\max})$. Recall that $\delta \leq \varepsilon \leq \gamma\delta$. We obtain that

$$\begin{aligned}
f(x_{\max}) \leq \frac{f(y^k)}{1-\varepsilon} &\leq \frac{(1+\varepsilon)^k}{1-\varepsilon} f(x_{\min}) \\
&\leq \frac{(1+\varepsilon)^{\log_{1+\delta}(x_{\max}/x_{\min})+1}}{1-\varepsilon} f(x_{\min}) \\
&\leq \frac{(1+\varepsilon)^{\log_{1+\delta}(4\lambda)+1}}{1-\varepsilon} f(x_{\min}) \\
&= \frac{1+\varepsilon}{1-\varepsilon} (4\lambda)^{\log_{1+\delta}(1+\varepsilon)} f(x_{\min}) \\
&\leq \frac{1+\varepsilon}{1-\varepsilon} (4\lambda)^{\log_{1+\delta}(1+\gamma\delta)} f(x_{\min}) \\
&= (1/\delta)^{O(1)} f(x_{\min}),
\end{aligned}$$

where the last expression follows since $\log_{1+\delta}(1+\gamma\delta) = \ln(1+\gamma\delta)/\ln(1+\delta) \leq \gamma\delta/\ln(1+\delta) = O(\gamma) = O(1)$ (for δ small enough), and since $\lambda = O(1/\delta)$. We conclude that

$$\max_{c \in Q} \bar{f}_c \leq (f(x_{\max}))/(\varepsilon f(x_{\min})) + 1 \leq (1/\varepsilon)(1/\delta)^{O(1)} + 1 = (1/\delta)^{O(1)}.$$

For the case in which $x_{\max} \leq x_{\min}$ we define the sequence $y^k := (1-\delta)^k x_{\min}$. The rest of the proof is analogous and the details are left to the reader. □

As we now must consider the objective function, we cannot simply apply Theorem 1 to [cost-conf-ILP]. However, we can prove a slightly weaker version by decomposing the ILP in several smaller ones and applying the theorem to each of them.

Theorem 16. *If [cost-conf-IP] is feasible, then there exists an optimal solution x satisfying:*

1. $\sum_{c \in Q_c} x_c \in O((1/\delta^3)\log^2(1/\delta))$, *and*
2. $|\operatorname{supp}(x) \cap Q_s| \in O((1/\delta)\log^2(1/\delta))$.

Proof. Notice that the load of each configuration $\pi \cdot c$ is a multiple of L/λ^2, and thus $\pi \cdot c \in \{L/\lambda, L/\lambda + L/(\lambda^2), \ldots, 4L\}$. We classify the configurations according to their loads, $Q^\ell := \{c \in Q : \pi \cdot c = L/\lambda + \ell \cdot L/(\lambda^2)\}$, for $\ell \in \{0, \ldots, 4\lambda^2 - \lambda\}$. Let x^* be an optimal solution of [cost-conf-IP]. Then we can considered an ILP for each load value ℓ:

$$[\text{conf-IP}]_\ell \quad \sum_{c\in Q^\ell} c\cdot x_c = \sum_{c\in Q^\ell} c\cdot x^*_c, \tag{4}$$

$$\sum_{c\in Q^\ell} x_c = \sum_{c\in Q^\ell} x^*_c, \tag{5}$$

$$x_c \in \mathbb{Z}_{\geq 0} \qquad \text{for all } c \in Q^\ell. \tag{6}$$

Scaling π by multiplying it by λ^2/L we obtain an integral vector (since π is an integer multiple of $L/(\lambda^2)$), we can apply Theorem 1 to each ILP $[\text{conf-IP}]_\ell$, which yields that there exists a thin solution x^ℓ. In particular the number of complex configurations in x^ℓ is $\sum_{c\in Q_c\cap Q^\ell} x^\ell_c \in O((1/\delta)\log^2(1/\delta))$. Since $\bar{f}_c$ depends only on the load of c, concatenating these solutions yields a solution $x' := (x^\ell)_\ell$ that is optimal for [cost-conf-IP], such that $\sum_{c\in Q_c} x'_c \in O((\lambda^2)\cdot(1/\delta)\log^2(1/\delta)) = O((1/\delta^3)\log^2(1/\delta))$. It remains to bound the number of simple configurations in the support. To this end, we consider the ILP restricted to simple configurations as follows:

$$[\text{cost-conf-IP}]_s \quad \min \sum_{c\in Q_s} \bar{f}_c\cdot x_c$$

$$\sum_{c\in Q_s} c\cdot x_c = b - \sum_{c\in Q_c} c\cdot x'_c, \tag{7}$$

$$\sum_{c\in Q_s} x_c = m - \sum_{c\in Q_c} x'_c, \tag{8}$$

$$x_c \in \mathbb{Z}_{\geq 0} \qquad \text{for all } c \in Q_s. \tag{9}$$

We apply the result of Eisenbrand and Shmonin [4] to this ILP. In its more general form, this result ensures the existence of a solution x'' with support of size $O(N(\log(N)+\Delta))$, where N is the number of restrictions and Δ is the encoding size of the largest coefficient appearing in the cost vector and restriction matrix. In our case $N = d+1 = O((1/\delta)\log(1/\delta))$, and $\Delta = O(\log(\max\{1/\delta, \max_{c\in Q}\bar{f}_c)\}) = O(\log(1/\delta))$ (Lemma 15). Thus $O(N(\log(N)+\Delta)) = O((1/\delta)\log^2(1/\delta))$. The theorem follows by concatenating $(x''_c)_{c\in Q_s}$ with $(x'_c)_{c\in Q_c}$. □

Finally, we use the structure given by the theorem to solve this ILP optimally.

Algorithm 17.

1. *For each processing time π_k, guess the number $b^c_k \leq b_k$ of jobs covered by complex configurations.*
2. *Guess the number m^c of machines that schedule jobs b^c.*
3. *Compute an optimal solution for instance with number of jobs b^c on m^c machines with a dynamic program.*
4. *Guess the support of simple configurations $\bar{Q}_s \subseteq Q_s$ used by the solution implied by Theorem 1, with $|\bar{Q}_s| \in O((1/\delta)\log^2(1/\delta))$.*
5. *Solve the ILP restricted to configurations in $\bar{Q}_s$:*

$$\min \sum_{c\in Q_s} \bar{f}_c\cdot x_c$$

$$\sum_{c\in \bar{Q}_s} c\cdot x_c = b - b^c,$$

$$\sum_{c\in \bar{Q}_s} x_c = m - m^c,$$

$$x_c \in \mathbb{Z}_{\geq 0} \qquad \text{for all } c \in \bar{Q}_s.$$

Lemma 18. *Algorithm 17 can be implemented with a running time of* $2^{O((1/\delta)\log^4(1/\delta))}\log(n)$.

Proof. In step 1, the algorithm guesses which jobs are processed on machines following a complex configurations. Since each configuration contains at most $O(1/\delta)$ jobs, there are at most $O((1/\delta^4)\log^2(1/\delta))$ jobs assigned to such machines. For each size $\pi_k \in \Pi$, we guess the number b_k^c of jobs of size π_k assigned to such machines. Hence, we can enumerate all possibilities for jobs assigned to complex machines in time $2^{O((1/\delta)\log^2(1/\varepsilon))}$. Similarly, we can guess the number of machines in step 2 since $m_c \in O((1/\delta^3)\log^2(1/\delta))$. For step 3 we use a simple dynamic program that goes over the machines storing a table $T(\ell, z_1, \ldots, z_d)$ that contains the minimum cost achieved over the first $\ell \leq m^c$ machines with $z_k \leq b_k^c$ jobs of size π_k. The number of entries the table is $O(m^c \prod_{k=1}^{d}(b_k^c+1))$. Computing $T(\ell, z_1, \ldots, z_d)$ can be done by checking all entries of the type $T(\ell-1, z_1', \ldots, z_d')$ for $z_k' \leq z_k$. Thus, the running time of the dynamic programm is $O(m^c[\prod_{k=1}^{d}(b_k^c+1)]^2)$. Since $b_k^c \in O((1/\delta^4)\log^2(1/\delta))$ for each k, recalling that $m^c \in O((1/\delta^3)\log^2(1/\delta))$, and that $d = |\Pi| \in O((1/\delta)\log(1/\delta))$, we obtain that step 3 can be implemented with $2^{O((1/\delta)\log^2(1/\delta))}$ running time.

In step 4, our algorithm guesses the support of the solution implied by Theorem 16. Let D be the bound implied by the third property of this theorem, so that $|\operatorname{supp}(x) \cap Q_s| \leq D$ and $D \in O((1/\delta)\log^2(1/\delta))$. Also, $|Q_s| \leq 2^{O(\log^2(1/\delta))}$ Then the guessing in step 4 needs to consider the following number of possibilities:

$$\sum_{i=0}^{D} \binom{|Q_s|}{i} \leq (D+1)|Q_s|^D \leq 2^{O((1/\delta)\log^4(1/\delta))}.$$

In step 5, the number of variables of the restricted ILP is $|\bar{Q}_s| = O((1/\delta)\log^2(1/\delta))$. Moreover, using Lemma 15 the size of the input can be bounded by $O((1/\delta^2)\log^3(1/\delta)\log(n))$. Running Kannan's algorithm [14] to solve the ILP takes time $2^{O((1/\delta)\log^3(1/\delta))}\log(n)$. Hence, the total running time of our algorithm can be bounded by $2^{O((1/\delta)\log^4(1/\delta))}\log(n)$. □

As in [2], the algorithm above can be easily adapted for objectives (II), (I') and (III') by suitably adapting the ILP. We leave the details to the reader. This suffices to conclude Theorem 19

Theorem 19. *Consider the scheduling problem on parallel machines with objective functions (I), (II) for f convex (respectively (I') and (II') for f concave). If f satisfies Condition 12 for $1/\delta = O(1/\varepsilon)$, then the problem admits an EPTAS with running time $2^{O((1/\varepsilon)\log^4(1/\varepsilon))} + O(n\log n)$.*

5 Minimum makespan scheduling on uniform machines

In this section we generalize our result for $P||C_{\max}$ to uniform machines. Consider a set of jobs $\mathcal{J}$ with processing times p_j and a set of m non-identical machines $\mathcal{M}$ where machine $i \in \mathcal{M}$ runs at speed s_i. If job j is executed on machine i the machine needs p_j/s_i time units to complete the job. The problem is to find an assignment $a : \mathcal{J} \to \mathcal{M}$ for the jobs to the machines that minimizes the makespan; $\max_i \sum_{j:a(j)=i} p_j/s_i$. The problem is denoted by $Q||C_{max}$. We suppose that $s_1 \geq s_2 \geq \ldots \geq s_m$. Jansen [11] found an efficient polynomial time approximation scheme (EPTAS) for this scheduling problem which has a running time of $2^{O(1/\varepsilon^2\log^3(1/\varepsilon))} + \operatorname{poly}(n)$. Here we show how to improve the running time and prove the main result of this section.

Theorem 20. *There is an EPTAS (a family of algorithms $\{A_\varepsilon : \varepsilon > 0\}$) which, given an instance I of $Q||C_{max}$ with n jobs and m machines and a positive number $\varepsilon > 0$, produces a schedule of makespan $A_\varepsilon(I) \leq (1+\varepsilon)\textsc{Opt}(I)$. The running time of A_ε is $2^{O(1/\varepsilon\log^4(1/\varepsilon))} + \operatorname{poly}(n)$.*

We follow the approach by Jansen [11], transforming the scheduling problem into a bin packing problem with different bin capacities, round the processing times and bin capacities, divide the bins into at most three groups and generate four different scenarios depending on the input instance.

First, we compute a 2-approximate solution using the algorithm by Gonzales et al. [6] of length $B(I) \leq 2\,\textsc{Opt}(I)$. Suppose that $\varepsilon < 1$; otherwise we can take the 2-approximate solution and

are done. Then we choose a value $\delta \in (0, \varepsilon)$ such that $1/\delta$ is integral (the exact value is specified later) and use a binary search within the interval $[B(I)/2, B(I)]$ (that contains $\text{OPT}(I)$). We use a standard dual approximation method that for each value T either computes an approximate schedule of length $T(1+a\delta)$ (where a is constant) or shows that there is no schedule of length T. Since $(\delta/2)B(I) \leq \delta\text{OPT}(I)$, we can find within $O(\log(1/\delta))$ iterations a value $T \leq \text{OPT}(I)(1+\delta)$ with a corresponding schedule of length at most $T(1+a\delta) \leq \text{OPT}(I)(1+\varepsilon)$, using $\delta \leq \varepsilon/(a+2)$ and $\varepsilon \leq 1$. Next, the scheduling problem is transformed into a bin packing problem with m bins and capacities $c_i = T \cdot s_i$, the processing times p_j are rounded to the next value $\bar{p}_j$ of the form $\delta(1+\delta)^{k_j}$ with $k_j \in \mathbb{Z}$ and the bin capacities are rounded to the next power of $(1+\delta)$. We call $\mathcal{B}$ the set of bins with rounded capacities.

Lemma 21 (Jansen [11]). *If there is a feasible packing of n jobs with processing times p_j into m bins with capacities c_i, then there is also a packing of n jobs with rounded processing times $\bar{p}_j = \delta(1+\delta)^{k_j} \leq (1+\delta)p_j$ into m bins with rounded bin capacities $c'_i = (1+\delta)^{\ell_i} \leq c_i(1+\delta)^2$ with $\ell_i \in \mathbb{Z}$.*

If the number m of bins is smaller than $K \in O(1/\delta \log(1/\delta))$, then we can use an approximation scheme by Jansen and Mastrolilli [12] to compute an $(1+\varepsilon)$-approximate solution to schedule n jobs on m unrelated machines (an even more general problem) within time $O(n) + 2^{O(m\log(m/\varepsilon))} = O(n) + 2^{O(1/\delta \log^2(1/\delta))} = O(n) + 2^{O(1/\varepsilon \log^2(1/\varepsilon))}$; using that $\delta \in O(\varepsilon)$. Suppose from now on that $K > O(1/\delta \log(1/\delta))$. Then, we divide the bins into at most three different bin groups. The first group $\mathcal{B}_1$ consists of the $K = O(1/\delta \log(1/\delta))$ largest bins. For some $\gamma \in \Theta(\varepsilon^2)$, the next group $\mathcal{B}_2$ consists either of all the remaining bins $\{b_{K+1}, \ldots, b_m\}$ if $c'_m > \gamma c'_K$ (and we have only two bin groups) or $\mathcal{B}_2$ contains the next G largest bins $\{b_{K+1}, \ldots, b_{K+G}\}$ where G is the smallest index such that capacity $c'_{K+G+1} \leq \gamma c'_K$. In the second case, $\mathcal{B}_3 = \{b_{K+G+1}, \ldots, b_m\}$. Let $c_{max}(\mathcal{B})$ and $c_{min}(\mathcal{B})$ be the largest and smallest bin capacity in $\mathcal{B}$. If $c_{max}(\mathcal{B})/c_{min}(\mathcal{B}) \leq C$ for some value C and $\mathcal{B}$ contains only rounded capacities $(1+\delta)^x$ with $x \in \mathbb{Z}$, then the number of different capacities in $\mathcal{B}$ is at most $O(1/\delta \log(C))$.

Lemma 22 (Jansen [11]). *If there is a solution for the original instance $(\mathcal{J}, \mathcal{M})$ of our scheduling problem with makespan T and corresponding bin sizes, then there is a feasible packing for instance $(\mathcal{J}, \mathcal{B}'_1 \cup \mathcal{B}_2 \cup \mathcal{B}_3)$ or instance $(\mathcal{J}, \mathcal{B}'_1 \cup \mathcal{B}_2)$ with rounded bin capacities $\bar{c}_i \leq c_i(1+\delta)^3$ and rounded processing times $\bar{p}_j \leq (1+\delta)p_j$. Here $\mathcal{B}'_1$ is the subset of $\mathcal{B}_1$ with bins of capacity larger than $\delta/(K-1)c_{max}(\mathcal{B}_1)$ and $\mathcal{B}_2$ has a constant number $O(1/\delta \log(1/\delta))$ of different bin capacities. In addition we have one of the following four scenarios:*

(1) *Two bin groups $\mathcal{B}'_1$ and $\mathcal{B}_2$ with a gap $c_{min}(\mathcal{B}'_1)/c_{max}(\mathcal{B}_2) \geq 1/\delta$.*

(2) *Two bin groups $\mathcal{B}_1$ and $\mathcal{B}_2$ with a constant number $O(1/\delta \log(1/\delta))$ of different bin capacities in $\mathcal{B}'_1 \cup \mathcal{B}_2$.*

(3) *Three bin groups $\mathcal{B}'_1, \mathcal{B}_2, \mathcal{B}_3$ with a gap $c_{min}(\mathcal{B}'_1)/c_{max}(\mathcal{B}_2) \geq 1/\delta$ and $c_{min}(\mathcal{B}_1)/c_{max}(\mathcal{B}_3) \geq 1/\gamma$.*

(4) *Three bin groups $\mathcal{B}'_1, \mathcal{B}_2, \mathcal{B}_3$ with a constant number $O(1/\delta \log(1/\delta))$ of different bin capacities in $\mathcal{B}'_1 \cup \mathcal{B}_2$ and a gap $c_{min}(\mathcal{B}_1)/c_{max}(\mathcal{B}_3) \geq 1/\gamma$.*

Notice that scenario 4 can be seen as a special case of scenario 3 by using $\mathcal{B}_2^{new} = \mathcal{B}_1 \cup \mathcal{B}_2$, $\mathcal{B}_1^{new} = \emptyset$, and $\mathcal{B}_3^{new} = \mathcal{B}_3$. The same modification works to show that scenario 2 is a special case of scenario 1. Finally, scenario 1 can be interpreted a special case of scenario 3, using $\mathcal{B}_3 = \emptyset$. Therefore, it is sufficient to improve the running time for scenario 3.

Scenario 3 will be solved using a mix of dynamic programming and mixed integer linear programming (MILP) techniques. In this approach we use only the larger bins in $\mathcal{B}'_1$ of $\mathcal{B}_1$ to execute jobs, but in the rounding step for scenario 3 afterwards we may also use the smaller bins in $\mathcal{B}_1$. Notice that a packing into a bin b_i with capacity $\bar{c}_i \leq c_i(1+\delta)^3$ corresponds to a schedule on machine i with total processing time at most $\bar{c}_i/s_i \leq T(1+\delta)^3$. For $T \leq \text{OPT}(I)(1+\delta)$ this gives us a schedule of length at most $\text{OPT}(I)(1+\delta)^4$. If there a feasible schedule with makespan T, then the total processing time of the instance is $\sum_{j\in\mathcal{J}} p_j \leq \sum_{i=1}^m \bar{c}_i$. If this inequality does

not hold, then we discard the choice with makespan T. Otherwise, we can eliminate the set $\mathcal{J}_{\text{tiny}}$ of tiny jobs with processing time $\leq \delta\bar{c}_m$ and pack them greedily at the end of the algorithm into the enlarged bins of size $\bar{c}_i(1+\delta)$. Hence, in what follows we assume that $\mathcal{J}_{\text{tiny}}$ is empty.

5.1 Solution for the instance $(\mathcal{J}, \mathcal{B}'_1 \cup \mathcal{B}_2 \cup \mathcal{B}_3)$

In this subsection we consider scenarios 3 above with three bin groups. First, we preassign all huge jobs with processing time $> \delta\bar{c}_{K'}$ into the first K' machines. Since $\bar{c}_{K+1} = c_{max}(\mathcal{B}_2) \leq \delta c_{min}(\mathcal{B}'_1) = \delta\bar{c}_{K'}$, the huge jobs fit only on the first K' bins. The number of huge jobs can be bounded by $O(Kc_{max}(\mathcal{B}_1)/(\delta c_{K'})) = O(1/\delta^4 \log^2(1/\delta))$. If there are more huge jobs in the instance, then there is no packing into $\mathcal{B}'_1 \cup \mathcal{B}_2$ and we are done. Furthermore, the number of machines $K' \in O(1/\delta \log(1/\delta))$ is constant. Again, we can use the approximation scheme by Jansen and Mastrolilli that computes an $(1+\delta)$-approximate schedule for N jobs on M machines which runs in $O(N) + 2^{O(M \log(M/\delta))}$. For $M \in O(1/\delta \log(1/\delta))$ and $N \in O(1/\delta^4 \log^2(1/\delta))$ this gives a running time $O(1/\delta^4 \log^2(1/\delta)) + 2^{O(1/\delta \log^2(1/\delta))} = 2^{O(1/\delta \log^2(1/\delta))}$ to obtain a feasible packing with bin sizes $\bar{c}_i(1+\delta)$ or schedule of length $\leq T(1+\delta)^4$, if one exists. If there is no feasible packing for the huge jobs, then there is no schedule with makespan T and we have to increase T in the binary search. In the other case we set up a MILP.

After the assignment of the huge jobs, we have a free area S_0 in $\mathcal{B}_1$ for the remaining jobs with processing time $\bar{p}_j \leq \delta\bar{c}_{K'}$. The different bin capacities in $\mathcal{B}_2$ and $\mathcal{B}_3$ are denoted by $\bar{c}(1) > \ldots > \bar{c}(L)$ and $\bar{c}(L+1) > \ldots > \bar{c}(L+N)$, respectively. Let m_ℓ be the number of bins of size $\bar{c}(\ell)$ for $\ell = 1, \ldots, L+N$. The m_ℓ machines of the same speed form a block B_ℓ of bins with the same capacity $\bar{c}(\ell)$. In addition, we have $n_1, \ldots, n_P$ jobs of size $\delta(1+\delta)^{k_j}$ and suppose that the first $P' \leq P$ job sizes are larger than $\bar{c}_{K+1} = \bar{c}(1)$.

In the MILP we use $C_1^{(\ell)}, \ldots, C_{h_\ell}^{(\ell)}$ as configurations or multisets with numbers $\delta(1+\delta)^{k_j} \in [\delta\bar{c}(\ell), \bar{c}(\ell)]$ (these are large processing times corresponding to block B_ℓ), where the total sum $size(C_i^{(\ell)}) = \sum_j a(k_j, C_i^{(\ell)})\delta(1+\delta)^{k_j}$ is bounded by $\bar{c}(\ell)$. Here $a(k_j, C_i^{(\ell)})$ is the number of occurrences of number $\delta(1+\delta)^{k_j}$ in configuration $C_i^{(\ell)}$. In the MILP below, we use integral and fractional variables $x_i^{(\ell)}$ to indicate number of machines that are scheduled according to configuration $C_i^{(\ell)}$. In addition, we use fractional variables $y_{j,\ell}$ to indicate the number of jobs of size $\delta(1+\delta)^{k_j}$ placed as small ones in block B_ℓ; i.e. $\delta(1+\delta)^{k_j} < \delta\bar{c}(\ell)$. For each job size $\delta(1+\delta)^{k_j} \leq \bar{c}(1)$, let a_j be the smallest index in $\{1, \ldots, L+N\}$ such that $\delta(1+\delta)^{k_j} \geq \delta\bar{c}(a_j)$. If there is no such index, we have a tiny processing time $\delta(1+\delta)^{k_j} < \delta\bar{c}(L+N)$. Notice that the first P' job sizes are within $(\bar{c}(1), \delta c_{K'}]$. These jobs do not fit into $\mathcal{B}_2 \cup \mathcal{B}_3$. Therefore, for these job sizes we use only one variable $y_{j,0} = n_j$ and set $a_j = 0$.

$$
\begin{array}{lll}
\sum_i x_i^{(\ell)} & \leq m_\ell & \text{for} \quad \ell = 1, \ldots, L+N, \\
\sum_{\ell,i} a(k_j, C_i^{(\ell)}) x_i^{(\ell)} + \sum_{\ell=0}^{a_j-1} y_{j,\ell} & = n_j & \text{for} \quad j = P'+1, \ldots, P, \\
\sum_i size(C_i^{(\ell)}) x_i^{(\ell)} + \sum_{j:\ell<a_j} y_{j,\ell}\delta(1+\delta)^{k_j} & \leq m_\ell \bar{c}(\ell) & \text{for} \quad \ell = 1, \ldots, L+N, \\
\sum_{j=1}^{P} y_{j,0}\delta(1+\delta)^{k_j} & \leq S_0, &
\end{array}
$$

$$
\begin{array}{ll}
x_i^{(\ell)} \text{ integral} \geq 0 & \text{for } \ell = 1, \ldots, L \text{ and } i = 1, \ldots, h_\ell, \\
x_i^{(\ell)} \geq 0 & \text{for } \ell = L+1, \ldots, L+N \text{ and } i = 1, \ldots, h_\ell \\
y_{j,0} = n_j & \text{for } j = 1, \ldots, P', \\
y_{j,\ell} \text{ integral} \geq 0 & \text{for } j = P'+1, \ldots, P \text{ and } \ell = 0, \ldots, a_j - 1.
\end{array}
$$

In the MILP above, we use integral variables for configurations in the blocks of group $\mathcal{B}_2$ and fractional variables for the configurations in blocks of $\mathcal{B}_3$. Each feasible packing for the jobs into the bins corresponds to a feasible solution of the MILP. The total number of variables is $O(n^2) + O(n)2^{O(1/\delta \log(1/\delta))}$, the number of integral variables is at most $2^{O(1/\delta \log(1/\delta))}$, and the number of constraints (not counting the non-negativity constraints) is at most $O(n)$. The previous approach to solve the scheduling problem and the underlying MILP had a running time of $2^{O(1/\delta^2 \log^3(1/\delta))} + \text{poly}(n)$. In order to use an approach similar to the scheduling on identical

machines, each large size $\delta(1+\delta)^{k_j} \in C_i^{(\ell)}$ is rounded up to the next multiple of $\delta^2\bar{c}(\ell)$. This enlarges the size of each configuration $C_i^{(\ell)}$ from $size(C_i^{(\ell)})$ to at most $size(C_i^{(\ell)}) + \delta\bar{c}(\ell)$ and the corresponding bin size from $\bar{c}(\ell)$ to $(1+\delta)\bar{c}(\ell)$.

Let $\bar{C}_1^{(\ell)}, \ldots, \bar{C}_{\bar{h}_\ell}^{(\ell)}$ be the configurations of size at most $(1+\delta)\bar{c}(\ell)$ with the rounded-up numbers $q(k_j,\ell)\delta^2\bar{c}(\ell)$ with $q(k_j,\ell) \in \mathbb{Z}^+$ and multiplicities $a(k_j, \bar{C}_i^{(\ell)})$. This rounding implies also that the rounded size $size(\bar{C}_i^{(\ell)})$ of a configuration is a multiple of $\delta^2\bar{c}(\ell)$. Each new rounded configuration $\bar{C}_i^{(\ell)}$ (with rounded-up numbers $q(k_j,\ell)\delta^2\bar{c}(\ell)$ and multiplicities $a(k_j, \bar{C}_i^{(\ell)})$) corresponds to an integral point inside the knapsack polytope $\mathcal{P}_\ell = \{C = (a(k_j, C)) : q \cdot C \leq 1/\delta^2 + 1/\delta\}$ such that $\sum_j q(k_j,\ell)a(k_j,\bar{C}_i^{(\ell)})\delta^2\bar{c}(\ell) = size(\bar{C}_i^{(\ell)}) \leq (1+\delta)\bar{c}(\ell)$ or, equivalently, $\sum_j q(k_j,\ell)a(k_j,\bar{C}_i^{(\ell)}) \leq 1/\delta^2 + 1/\delta \leq 2/\delta^2$. We consider now a modified MILP with configurations $\bar{C}_i^{(\ell)}$ and coefficients $a(k_j, \bar{C}_i^{(\ell)})$. Note that the total area of all configurations in B_ℓ can be bounded by $\sum_i size(\bar{C}_i^{(\ell)})x_i^{(\ell)} \leq \sum_i size(C_i^{(\ell)})x_i^{(\ell)} + \delta\bar{c}(\ell)\sum_i x_i^{(\ell)}$. This, together with the small jobs gives $\sum_i size(\bar{C}_i^{(\ell)})x_i^{(\ell)} + \sum_j \delta(1+\delta)^{k_j}y_{j,\ell} \leq \sum_i size(C_i^{(\ell)})x_i^{(\ell)} + \delta m_\ell\bar{c}(\ell) + \sum_j \delta(1+\delta)^{k_j}y_{j,\ell} \leq m_\ell\bar{c}(\ell)(1+\delta)$; i.e. the total area is increased by at most a multiplicative factor of $(1+\delta)$. Since the total area of all jobs within one block is increased by this rounding, we use the following new constraints in the modified MILP:

$$\sum_i size(\bar{C}_i^{(\ell)})x_i^{(\ell)} + \sum_j y_{j,\ell}\delta(1+\delta)^{k_j} \leq m_\ell\bar{c}(\ell)(1+\delta) \qquad \text{for } \ell = 1, \ldots, L.$$

Next, we divide the coefficients in the L area constraints above by $\delta^2\bar{c}(\ell)$. Then the coefficients of the $x_i^{(\ell)}$ variables are now $size(\bar{C}_i^{(\ell)})/(\delta^2\bar{c}(\ell)) = a_{i,\ell}\delta^2\bar{c}(\ell)/(\delta^2\bar{c}(\ell)) = a_{i,\ell} \in \{1/\delta, \ldots, 1/\delta^2 + 1/\delta\}$. Using the assumption that $1/\delta$ is integral, all coefficients of the variables are integral and bounded by $2/\delta^2$. Notice that increasing the capacities of all bins and dividing all coefficients as above, implies also a feasible solution of the modified MILP. Let us study a feasible solution of the modified MILP. To reduce the number of integral configuration variables in the MILP, we consider the following ILP that uses only the integral $x_i^{(\ell)}$ variables within bin group $\mathcal{B}_2$:

$$\sum_i x_i^{(\ell)} = \bar{m}_\ell \qquad \text{for } \ell = 1, \ldots, L, \tag{10}$$

$$\sum_{\ell,i} a(k_j, \bar{C}_i^{(\ell)})x_i^{(\ell)} = \bar{n}_j \qquad \text{for } j \in P(\mathcal{B}_2), \tag{11}$$

$$\sum_i \frac{size(\bar{C}_i^{(\ell)})}{\delta^2\bar{c}(\ell)}x_i^{(\ell)} = Area(\ell, large) \qquad \text{for } \ell = 1, \ldots, L, \tag{12}$$

$$x_i^{(\ell)} \text{ integral} \geq 0 \qquad \text{for } i = 1, \ldots, \bar{h}_\ell, \ell = 1, \ldots, L. \tag{13}$$

where the values $\bar{m}_\ell$, $\bar{n}_j$, and $Area(\ell, large)$ are given by a feasible solution of the modified MILP. Here $P(\mathcal{B}_2)$ is the set of all indices of large job sizes corresponding to blocks $B_\ell \in \mathcal{B}_2$; i.e. $P(\mathcal{B}_2) = \{j : \delta(1+\delta)^{k_j} \in (\delta\bar{c}(L), \bar{c}(1)]\}$. The cardinality of $P(\mathcal{B}_2)$ and the value L can be bounded by $O(1/\delta\log(1/\delta))$. All the coefficients above of the variables are bounded by $O(1/\delta^2)$.

The support of a configuration $\bar{C}_i^{(\ell)}$ is the number of values $a(k_j, \bar{C}_i^{(\ell)}) > 0$; i.e. $\text{supp}(\bar{C}_i^{(\ell)}) = |\{j : a(k_j, \bar{C}_i^{(\ell)}) > 0\}|$. In our case $\text{supp}(\bar{C}_i^{(\ell)})) \leq O(1/\delta\log(1/\delta))$. A configuration $\bar{C}_i^{(\ell)}$ is called simple, if $|\text{supp}(\bar{C}_i^{(\ell)})| \leq \log(1/\delta^2 + 1/\delta + 1)$ Otherwise, we call a configuration $\bar{C}_i^{(\ell)}$ complex. Using the result by Eisenbrand and Shmonin, we can find a feasible solution of the ILP above (if there is a feasible solution of the modified MILP) with at most $O(1/\delta\log^2(1/\delta))$ many variables $x_i^{(\ell)} > 0$; i.e. $|\text{supp}(x)| \leq O(1/\delta\log^2(1/\delta))$ where $x = (x_i^{(\ell)})$. We can generalize our result in Theorem 1 to our ILP above.

Lemma 23. *Assume that the ILP defined by* (10)-(13) *is feasible and let S denote the set of all simple configurations. Then there exists a feasible solution x' such that:*

(1) *If $x'^{(\ell)}_i > 1$ then the configuration $\bar{C}_i^{(\ell)}$ is simple.*

(2) *The support of x' satisfies* $|\operatorname{supp}(x') \cap S| \in O(1/\delta \log^2(1/\delta))$.

(3) *The support of x' satisfies* $|\operatorname{supp}(x') \setminus S| \in O(1/\delta^2 \log^3(1/\delta))$.

Proof. As stated above, the set of configurations $\bar{C}_1^{(\ell)}, \ldots, \bar{C}_{\bar{h}_\ell}^{(\ell)}$ equals the set of integral points Q_ℓ inside the knapsack polytope $\mathcal{P}_\ell = \{C = (q(k_j, C)) : q \cdot C \leq 1/\delta^2 + 1/\delta\}$ Let $\bar{x} = (\bar{x}^{(\ell)})_{\ell=1}^{L}$ be a solution to (10)-(13) where $\bar{x}^{(\ell)}$ corresponds to the variables defining the solution for block B_ℓ. We consider a family of ILPs defined for each $\ell = 1, \ldots, L$.

$$
\begin{aligned}
[\text{conf-IP}]_\ell \quad & \sum_{c \in Q_\ell} c \cdot x_c = \sum_{c \in Q_\ell} c \cdot \bar{x}_c^{(\ell)}, \\
& \sum_{c \in Q_\ell} x_c = \bar{m}_\ell, \\
& x_c \in \mathbb{Z}_{\geq 0} \qquad\qquad \text{for all } c \in Q_\ell.
\end{aligned}
$$

Using Theorem 1 for each [conf-IP]$_\ell$, we obtain new solution $\hat{x}^{(\ell)}$, where each complex configuration is used at most once and $\operatorname{supp}(\hat{x}^{(\ell)}) \in O(1/\delta \log^2(1/\delta))$. Then we define a new solution $\hat{x}$ of ILP (10)-(13) defined as $(\hat{x}^{(\ell)})_\ell$. In $\hat{x}$ every complex configuration is used at most once and $|\operatorname{supp}(\hat{x})| \leq L \cdot 2(\bar{d}+1)\log(4(\bar{d}+1)\bar{T}) \in O(1/\delta^2 \log^3(1/\delta))$, where $\bar{d} \leq |P(\mathcal{B}_2)| \in O(1/\delta \log 1/\delta)$ and $\bar{T} \in O(1/\delta^2)$. Note that Equation (12) of the above ILP holds for the new solution $\hat{x}$ as the set of jobs covered inside a block does not change and hence

$$
\sum_i \frac{size(\bar{C}_i^{(\ell)})}{\delta^2 \bar{c}(\ell)} \hat{x}_i^{(\ell)} = Area(\ell, large) = \sum_i \frac{size(\bar{C}_i^{(\ell)})}{\delta^2 \bar{c}(\ell)} \bar{x}_i^{(\ell)}. \qquad \square
$$

Finally, consider the ILP (10)-(13) and fix each variable x_i^ℓ, for $\bar{C}_i^{(\ell)}$ a complex configuration, to the value $\hat{x}_i^\ell$ (and thus the resulting ILP has variables only for simple configurations). Now we can apply the result of Eisenbrand and Shmonin [4] to this ILP. This ensures that any ILP of the form $\{z \in \mathbb{Z}_{\geq 0} : Az = h\}$ admits a solution with support of size $O(N(\log(N) + \Delta))$, where N is the number of rows of A and Δ is the largest encoding size of an entry of A. Recalling that $\frac{size(\bar{C}_i^{(\ell)})}{\delta^2 \bar{c}(\ell)} \in O(1/\delta^2)$, we can apply this result to our case, which yields a solution whose support contains at most $O(1/\delta \log^2(1/\delta))$ simple configurations. Hence, we obtain a solution satisfying all properties of the statement of the theorem.

Algorithm 24.

1. *For each job size, guess the number of jobs $v_j \leq \bar{n}_j$ covered by complex configurations.*
2. *For each bin size, guess the number of machines $w_j \leq \bar{m}_j$ used to schedule the set of jobs covered by complex configurations.*
3. *For each block ℓ in $\mathcal{B}_2$, guess the support of simple configurations $\bar{Q}_s^{(\ell)} \subseteq Q_s^{(\ell)}$ used by a thin solution, with $\sum_{\ell=1}^{L} |\bar{Q}_s^{(\ell)}| \leq 4(d+1)\log(4(d+1)\bar{T}) \in O((1/\varepsilon)\log^2(1/\varepsilon))$.*
4. *Solve the reduced modified MILP, where the integral variables x_i^ℓ are restricted to simple configurations.*

Lemma 25. *Algorithm 24 can be implemented with a running time of $2^{O((1/\varepsilon)\log^4(1/\varepsilon))} poly(n)$.*

Proof. As in the case of identical machines, our algorithm guesses in step 1 the complex configurations and the corresponding jobs. Since the number M of complex configurations within $\mathcal{B}_2$ ist at most $O(1/\delta^2 \log^3(1/\delta))$ and there are at most $O(1/\delta)$ many large job per configuration, the total number N of jobs within the complex configurations is at most $O(1/\delta^3 \log^3(1/\delta))$.

To obtain a schedule for the guessed jobs, notice that the number of large job sizes $|P(\mathcal{B}_2)| \leq O(1/\delta \log(1/\delta))$. We guess now a vector $v = (v_j)$ with possible job sizes that are covered by the complex configurations. The total number of these vectors is $(N+1)^{|P(\mathcal{B}_2)|} \leq (1/\delta^3 \log^3(1/\delta))^{O(1/\delta \log(1/\delta))} =$

$2^{O(1/\delta \log^2(1/\delta))}$. In addition, we guess a vector $w = (w_\ell)$ with the numbers w_ℓ of complex configurations in the block groups B_ℓ. The number of choices here is at most $(M+1)^L \leq (1/\delta^2 \log^3(1/\delta))^{O(1/\delta \log(1/\delta))} = 2^{O(1/\delta \log^2(1/\delta))}$. For each guess v, w we run a dynamic program to test whether the number of job sizes, stored in v, fit on the corresponding machines in the blocks B_ℓ, given by vector w. To do this, we run over the machines and store after ℓ machines, for $\ell = 1, \ldots, M$, the set of all feasible vectors with job sizes that can be packed into the first ℓ machines. This dynamic program runs in time $M 2^{O(1/\delta \log^2(1/\delta))} = 2^{O(1/\delta \log^2(1/\delta))}$. For each feasible choice of v, w we compute the reduced MILP by $\hat{m}_\ell = m_\ell - w_\ell$ and $\hat{n}_j = n_j - v_j$ and guess the support of a feasible solution x in the MILP; i.e. the simple configurations in $\mathcal{B}_2$ with value $x_i^{(\ell)} > 0$. The total number of simple configurations $Q_s^{(\ell)}$ in one bin block can be bounded, using observation 7, by $2^{O(\log^2(1/\delta))}$. Therefore, the total number of simple configurations in $\mathcal{B}_2$ is $\sum_{\ell=1}^{L} |Q_s^{(\ell)}| \leq L \cdot 2^{O(\log^2(1/\delta))} = 2^{O(\log^2(1/\delta))}$. This implies that the number of choices for the support of x is at most

$$\binom{\sum_\ell |Q_s^{(\ell)}|}{O(1/\delta \log^2(1/\delta))} = \binom{2^{O(\log^2(1/\delta))}}{O(1/\delta \log^2(1/\delta))} = 2^{O(1/\delta \log^4(1/\delta))}.$$

For each choice we solve a reduced MILP with $d = O(1/\delta \log^2(1/\delta))$ integral variables (step 4). The total size s of the MILP can be bounded by $s \leq poly(n, 1/\delta) + n \log(n) 2^{O(1/\delta \log(1/\delta))}$. Using the algorithm by Kannan with runnning time $d^{O(d)} poly(s)$ for an MILP with d variables and size s, we obtain a running time to solve one MILP in time $2^{O(1/\delta \log^3(1/\delta))} poly(n)$; using $poly(s) \leq poly(n) 2^{O(1/\delta \log(1/\delta))}$. Running over all vectors v, w and all guesses for the simple configurations, we obtain a running time of $2^{O(1/\delta \log^4(1/\delta))} + poly(n)$. □

The rounding of the fractional variables in the MILP solutions and the packing of the items accordingly works as in [11] and can be done in time $2^{O(1/\delta \log(1/\delta))} poly(n)$. Therefore, the overall running time of the entire algorithm can be bounded by $2^{O(1/\delta \log^4(1/\delta))} + poly(n) = 2^{O(1/\varepsilon \log^4(1/\varepsilon))} + poly(n)$; using that $\delta \in O(\varepsilon)$.

In order to calculate the length of the computed schedule and to specify δ, we use the following result:

Lemma 26. *[11] If there is a feasible solution of an MILP instance with bin capacities $\bar{c}(\ell)$ for blocks $B_\ell \in \mathcal{B}_2 \cup \mathcal{B}_3$ and capacities $\bar{c}_i$ for the K largest bins in $\mathcal{B}_1$, then the entire job set $\mathcal{J}$ can be packed into bins with capacities $\bar{c}(\ell)(1+2\delta)^2$ for blocks $B_\ell \in \mathcal{B}_2 \cup \mathcal{B}_3$ and enlarged capacities $\bar{c}_i(1+3\delta)^2$ for the first K bins.*

Note that the result above is constructive, too. This means that there is also an algorithm that computes a corresponding packing [11]. Using $\bar{c}_i \leq c_i(1+\delta)^3$ and $T \leq (1+\delta)OPT$ and the lemma above, we can bound the schedule length. If there is a schedule with length at most T and with corresponding bin sizes $c_i = Ts_i$, then the lemma above implies a packing into bins of size $c_i(1+3\delta)^3(1+3\delta)^2$ and a corresponding schedule length $\leq T(1+\delta)^3(1+3\delta)^2 \leq OPT(1+\delta)^4(1+3\delta)^2 \leq OPT(1+16\delta) \leq OPT(1+\varepsilon)$ for $\delta \leq \varepsilon/16$ and $\varepsilon \leq 1$. Using $\delta = \frac{1}{\lceil 16/\varepsilon \rceil}$, we obtain $\delta \leq \varepsilon/16$, $\delta \geq \varepsilon/17$, and that $1/\delta = \lceil 16/\varepsilon \rceil$ is integral. This concludes the proof for Theorem 20.

References

[1] N. Alon, Y. Azar, G. Woeginger, and T. Yadid. Approximation schemes for scheduling. In *Proceedings of the 8th Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '97, pages 493–500, 1997.

[2] N. Alon, Y. Azar, G. J. Woeginger, and T. Yadid. Approximation schemes for scheduling on parallel machines. *Journal of Scheduling*, 1:55–66, 1998.

[3] L. Chen, K. Jansen, and G. Zhang. On the optimality of approximation schemes for the classical scheduling problem. In *Proceedings of the 25th Annual ACM-SIAM Symposium*

on Discrete Algorithms (SODA 2013), pages 657–668. Society for Industrial and Applied Mathematics, 2013.

[4] F. Eisenbrand and G. Shmonin. Carathéodory bounds for integer cones. *Operations Research Letters*, 34:564–568, 2006.

[5] M. X. Goemans and T. Rothvoß. Polynomiality for bin packing with a constant number of item types. In *Proceedings of the 25th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA '14)*, pages 830–839, 2014.

[6] T. F. Gonzalez, O. H. Ibarra, and S. Sahni. Bounds for LPT schedules on uniform processors. *SIAM Journal on Computing*, 6(1):155–166, 1977.

[7] R. L. Graham. Bounds for certain multiprocessing anomalies. *Bell System Technical Journal*, 45:1563–1581, 1966.

[8] R. L. Graham. Bounds on multiprocessing timing anomalies. *SIAM Journal on Applied Mathematics*, 17:416–429, 1969.

[9] D. Hochbaum, editor. *Approximation algorithms for NP-hard problems*. PWS Publishing Company, 1997.

[10] D. S. Hochbaum and D. B. Shmoys. Using dual approximation algorithms for scheduling problems: theoretical and practical results. *Journal of the ACM*, 34:144–162, 1987.

[11] K. Jansen. An EPTAS for scheduling jobs on uniform processors: Using an MILP relaxation with a constant number of integral variables. *SIAM Journal on Discrete Mathematics*, 24:457–485, 2010.

[12] K. Jansen and M. Mastrolilli. Scheduling unrelated parallel machines: linear programming strikes back. *University of Kiel, Technical Report*, (1004), 2010.

[13] K. Jansen and C. Robenek. Scheduling jobs on identical and uniform processors revisited. In R. Solis-Oba and G. Persiano, editors, *Approximation and Online Algorithms*, number 7164 in Lecture Notes in Computer Science, pages 109–122. 2011.

[14] R. Kannan. Minkowski's convex body theorem and integer programming. *Mathematics of Operations Research*, 12:415–440, 1987.

[15] H. W. Lenstra. Integer programming with a fixed number of variables. *Mathematics of Operations Research*, 8:538–548, 1983.

[16] J. Y-T. Leung. Bin packing with restricted piece sizes. *Information Processing Letters*, 31:145–149, 1989.